arXiv:1511.05526v2 [cs.RO] 1 Jul 2016

Learning Articulated Motion Models from Visual and Lingual Signals

Zhengyang Wu
Georgia Tech
Atlanta, GA 30332
zwu66@cc.gatech.edu

Mohit Bansal
TTI-Chicago
Chicago, IL 60637
mbansal@ttic.edu

Matthew R. Walter
TTI-Chicago
Chicago, IL 60637
mwalter@ttic.edu

***Abstract*— In order for robots to operate effectively in homes and workplaces, they must be able to manipulate the articulated objects common within environments built for and by humans. Previous work learns kinematic models that prescribe this manipulation from visual demonstrations. Lingual signals, such as natural language descriptions and instructions, offer a complementary means of conveying knowledge of such manipulation models and are suitable to a wide range of interactions (e.g., remote manipulation). In this paper, we present a multimodal learning framework that incorporates both visual and lingual information to estimate the structure and parameters that define kinematic models of articulated objects. The visual signal takes the form of an RGB-D image stream that opportunistically captures object motion in an unprepared scene. Accompanying natural language descriptions of the motion constitute the lingual signal. We present a probabilistic language model that uses word embeddings to associate lingual verbs with their corresponding kinematic structures. By exploiting the complementary nature of the visual and lingual input, our method infers correct kinematic structures for various multiple-part objects on which the previous state-of-the-art, visual-only system fails. We evaluate our multimodal learning framework on a dataset comprised of a variety of household objects, and demonstrate a 36% improvement in model accuracy over the vision-only baseline.**

I. Introduction

As robots move off factory floors and into our homes and workplaces, they face the challenge of interacting with the articulated objects frequently found in environments built by and for humans (e.g., drawers, ovens, refrigerators, and faucets). Typically, this interaction is predefined in the form of a manipulation policy that must be (manually) specified for each object that the robot is expected to interact with. In an effort to improve efficiency and generalizability, recent work employs visual demonstrations to learn representations that describe the motion of these parts in the form of kinematic models that express the rotational, prismatic, and rigid relationships between object parts [1–4]. These structured models, which constrain the manifold on which the object's motion lies, allow for manipulation policies that are more efficient and deliberate. However, such visual cues may be too time-consuming to provide or may not be readily available, such as in the case of a disaster relief scenario in which a user is remotely commanding a robot over a bandwidth-limited channel. Further, reliance solely on vision makes these methods sensitive to common errors in object segmentation and tracking that occur as a result of clutter, occlusions, and a lack of visual features. Consequently, most existing systems require scenes to be free of distractors and that object parts be labeled with fiducial markers.

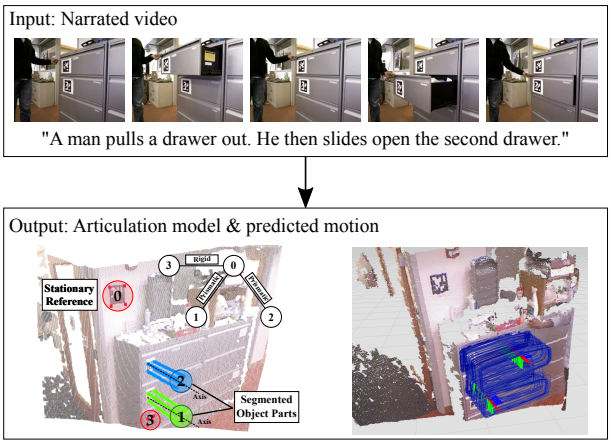

Fig. 1. Our framework learns the kinematic model that governs the motion of articulated objects (lower-left) from narrated RGB-D videos. The method can then use this learned model to subsequently predict the motion of an object's parts (lower-right).

Lingual input in the form of natural language descriptions and instructions offer a flexible, bandwidth-efficient medium that humans can readily use to convey knowledge of an object's operation. Such lingual descriptions of an articulated motion also provide a source of information that is complementary to visual input. Thus, these descriptions can be used to overcome some of the limitations of using visual-only observations, e.g., by providing cues regarding the number of parts that comprise the object or the motion type (e.g., rotational) between a pair of parts. In this work, we present a multimodal learning framework that estimates the kinematic structure and parameters of complex multi-part objects using both visual and lingual input, and performs substantially better than visual-only systems.

Our effort is inspired by the recent attention that has been paid to the joint use of vision and language as complementary signals for machine perception [5–22]. Much of the work in multimodal learning considers the problems of image caption generation and visual coreference resolution. Instead, we leverage the joint advantages of these two modalities in order to estimate the structure and parameters that define kinematic models of complex, multi-part objects such as doors, desks, chairs, and appliances.

Our multimodal learning framework first extracts noisy observations of the object parts and their motion separately from the visual and lingual signals. It then fuses these observations to learn a probabilistic model over the kinematic structure and model parameters that best explain the motion observed in the visual and lingual streams. Integral to this process is an appropriate means of representing the ambiguous nature of observations gleaned from natural language descriptions. We propose two probability models that capture this uncertainty based upon the similarity between the natural language text and a representative reference word set (for each model type) in a word embedding space. The first takes the form of a hard assignment of verbs in the description to the nearest kinematic model type (e.g., rotational or prismatic) in the embedding space. The second takes the form of a soft assignment, representing the likelihood of the lingual observations in terms of the similarity between the input text and the reference embeddings.

Our contributions include a multimodal approach to learning kinematic models from visual and lingual signals, the exploration of different language grounding methods to align action verbs and kinematic models, and the examination of various language priors in our learning framework. We evaluate our method on a dataset of video-text pairs demonstrating the motion of common household objects, and achieve notable improvement over the previous state-of-the-art, which only uses visual information. The word embedding-based soft and hard language models yield improvements of 21% and 36%, respectively, demonstrating the promise of a multimodal learning framework that exploits both visual and lingual information.

II. Related Work

Our goal is to enable robots to learn kinematic models with minimal supervision from human demonstrations. This requires solutions that can mitigate the complexity and clutter typical of human-occupied environments, without the need for additional infrastructure (e.g., visual fiducials).

Recent work considers the problem of learning articulated models based upon visual observations of demonstrated motion. Several methods formulate this problem as bundle adjustment, using structure-from-motion methods to first segment an articulated object into its compositional parts and to then estimate the parameters of the rotational and prismatic degrees-of-freedom that describe inter-part motion [23, 2]. These methods are prone to erroneous estimates of the pose of the object's parts and of the inter-part models as a result of outliers in visual feature matching. Alternatively, Katz et al. [24] propose an active learning framework that allows a robot to interact with articulated objects to induce motion. This method operates in a deterministic manner, first assuming that each part-to-part motion is prismatic. Only when the residual error exceeds a threshold does it consider the alternative rotational model. Further, they estimate the models based upon interactive observations acquired in a structured environment free of clutter, with the object occupying a significant portion of the RGB-D sensor's field-of-view. Katz et al. [3] improve upon the complexity of this method while preserving the accuracy of the inferred models. This method is prone to over-fitting to the observed motion and may result in overly complex models to match the observations. Hausman et al. [25] similarly enable a robot to interact with the object and describe a probabilistic model that integrates observations of fiducials with manipulator feedback. Meanwhile, Sturm et al. [1] propose a probabilistic approach that simultaneously reasons over the likelihood of observations while accounting for the learned model complexity. Their method requires that the number of parts that compose the object be known in advance and that fiducials be placed on each part to enable the visual observation of motion. More recently, Pillai et al. [4] propose an extension to this work that uses novel vision-based motion segmentation and tracking that enables model learning without prior knowledge of the number of parts or the placement of fiducial markers. Our approach builds upon this method with the addition of natural language descriptions of motion as an additional observation mode in a multimodal learning framework.

Meanwhile, recent work in the natural language processing community has focused on the role of language as a means of commanding [26–30, 22] and sharing spatial information [31–33] with robots. We use language for the novel and more complex task of learning object articulation in terms of kinematic motion models. Meanwhile, other methods have similarly used visual and lingual cues in a multimodal learning framework for such tasks as image and video caption synthesis [5–7, 10–21], visual coreference resolution [8, 9], visual question-answering [34], and understanding cooking videos paired with recipes [35]. Our work shares similar goals, particularly in the context of action inference based on joint visual-lingual cues.

III. Multimodal Learning Framework

Given an RGB-D video paired with the corresponding natural language description (alternatively, an instruction or caption) of an articulated object's motion, our goal is to infer the structure and parameters of the object's kinematic model. Adopting the formulation proposed by Sturm et al. [1], we represent this model as a graph, where each vertex denotes a different part of the object (or the stationary background) and edges denote the existence of constrained motion (e.g., a linkage) between two parts (Fig. 1). More formally, we estimate a *kinematic graph* $G = (V_G, E_G)$ that consists of vertices V_G for each object part and edges $E_G \subset V_G \times V_G$ between parts whose relative motion is kinematically constrained. Associated with each edge $(ij) \in E_G$ is its kinematic type $M_{ij} \in \{\text{rotational}, \text{prismatic}, \text{rigid}\}$ as well as the corresponding parameters θ_{ij}, such as the axis of rotation and the range of motion (see Fig. 2, lower-right). We take as input visual D_v and lingual D_l observations of the type and parameters of the edges in the graph. Our method then uses this vision-language observation pair $D_z = \{D_v, D_l\}$ to infer the maximum a posteriori kinematic structure and

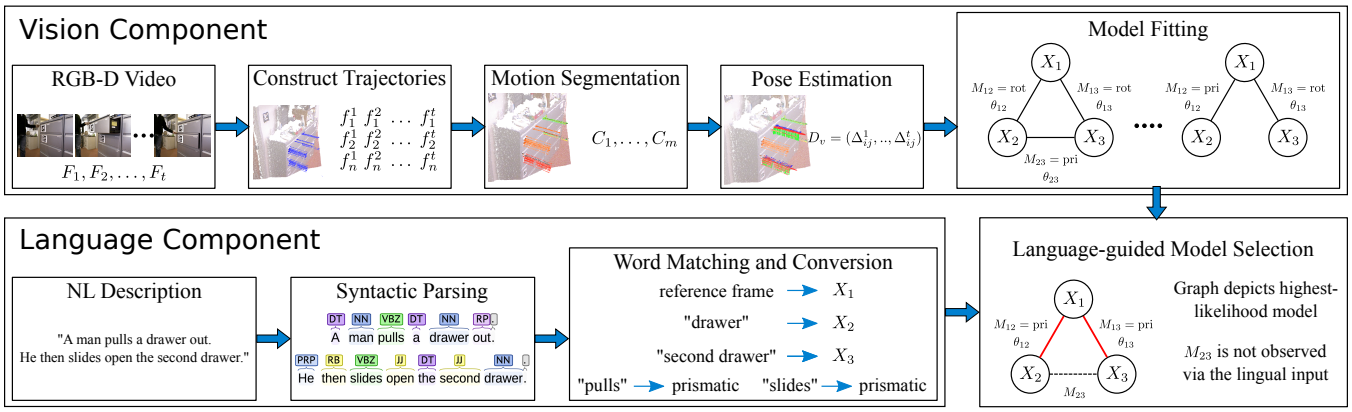

Fig. 2. Our multimodal articulation learning framework first identifies clusters of visual features that correspond to individual object parts. It then uses these feature trajectories to estimate the model parameters, assuming an initial estimate of the kinematic type associated with each edge in the graph. The method then uses natural language descriptions of the motion to estimate the kinematic type of each edge through a probabilistic language model.

model parameters that constitute the kinematic graph:

$$\hat{G} = \arg\max_{G} p(G|D_z) \tag{1a}$$

$$= \arg\max_{G} p(\{M_{ij}, \theta_{ij}|(ij) \in E_G\}|D_z) \tag{1b}$$

$$= \arg\max_{G} \prod_{(ij)\in E_G} p(M_{ij}, \theta_{ij}|D_z) \tag{1c}$$

Due to the complexity of joint inference, we adopt the procedure described by Sturm et al. [1] and use a two-step inference procedure that alternates between model parameter fitting and model structure selection steps (Fig. 2). In the first step, we assume a particular kinematic model type between each object i and j (e.g., prismatic), and then estimate the kinematic parameters based on the vision data (relative transformation between the two objects) and the assumed model type M_{ij}. We make one such assumption for each possible model type for each object pair.

In the model selection step, we then use the natural language description to infer the kinematic graph structure that best expresses the observation. While our previous work [4] provides visual observations of motion without the need for fiducials, it relies upon feature tracking and segmentation that can fail when the object parts lack texture (e.g., metal door handles) or when the scene is cluttered. Our system incorporates language as an additional, simpler, complementary observation of the motion, in order to improve the robustness and accuracy of model selection.

A. *Vision-guided Model Fitting*

Given the RGB-D video of the motion, we employ the vision pipeline of Pillai et al. [4] to arrive at a visual observation of the trajectory of each object part (Fig. 2). The method first identifies a set of 3D feature trajectories that correspond to different elements in the scene, including the object parts, background, and clutter (the distinction of which is not known a priori). These trajectories are then grouped to form rigid clusters according to the similarity of their relative motion in an effort to associate a cluster to each object part as well as to the background. Next, the method estimates the 6-DOF pose trajectory of each cluster (object part). We refer the reader to Pillai et al. [4] for specific details regarding the visual pipeline.

The 6-DOF pose trajectories constitute the visual observation of the motion D_v. Our framework uses these trajectories to estimate the parameters of a candidate kinematic model during the model fitting step. Specifically, we find the kinematic parameters that best explain the visual data given the assumed model

$$\hat{\theta}_{ij} = \arg\max_{\theta_{ij}} p(D_v|\hat{M}_{ij}, \theta_{ij}), \tag{2}$$

where $D_v = (\Delta_{ij}^1, ..., \Delta_{ij}^t), \forall (ij) \in E_G$ is the sequence of observed relative transformations between the poses of two object parts i and j, and $\hat{M}_{ij}$ is the current estimate of their model type. We perform this optimization over the joint kinematic structure defined by the edges in the graph [1].

B. *Language-guided Model Selection*

Methods that solely use visual input are sensitive to the effects of scene clutter and the lack of texture, which can result in erroneous estimates for the structure and parameters of the kinematic model [4]. We incorporate lingual observations into our framework in order to reduce errors that result from the failure in the visual pipeline, and also to add complementary observational information.

Specifically, we consider a natural language caption D_l that describes the motion observed in the video. Given this description, we infer the maximum a posteriori model type for each pair of object parts according to the caption. We require that the natural language descriptions adhere to a known grammar and that it narrate at least one motion present in the video[1] (otherwise, our method estimates the kinematic graph based solely on the visual observation). Note that we do not assume that valid captions provide an unambiguous description of the motion, but rather consider a distribution over the lingual observation, which provides

[1]As determined by the number of verbs.

Prismatic: *pull*, *push*, shift, *move*, *close*, remove, tug, yank, dislocate, extract, jerk, thrust, poke, prod, shove, displace, stretch, squeeze, fasten, draw, join, insert, embed, enter, exit, implant, inject, introduce, stick, admit, infuse, inlay, instill, place, set, penetrate, withdraw, intrude, slide.

Rotational: bend, yaw, turn, spin, whirl, *move*, *pull*, *push*, *close*, revolve, rotate, gyre, gyrate, pivot, swivel, twist, twirl, circle, roll, reel, wheel, round, wrench, screw, tighten, swing, cycle, bow, flex, wind, spiral, twine, loosen.

Fig. 3. Our manual dictionary of motion verbs for prismatic and rotational kinematic types. Words in italics are shared between the two dictionaries.

robustness to "noisy" captions. We employ the following procedure (bottom-left of Fig. 2) to convert a natural language description into a structured caption representation:

1) Perform word tokenization and part-of-speech tagging of the natural language description to obtain object nouns and action verbs.
2) Align the object nouns to the visual motion trajectories.
3) Classify the action verbs into kinematic model types (i.e., "prismatic," "rotational," or "rigid").

Next, we discuss each of these steps in detail.

1) Preprocessing Natural Language: Our system first extracts object- and motion-relevant cues from the natural language caption in the form of nouns that denote object parts and verbs that describe their motion. Nouns that refer to the agent (e.g., "man" or "person") are ignored. We use the Stanford CoreNLP pipeline, tokenizer, and POS-tagger to identify the nouns and verbs in the caption [36].

2) Matching Object Nouns with Trajectories: Given the set of nouns in the narration, we next seek to identify the corresponding object parts in the visual clusters. We enumerate the space of possible noun-cluster correspondences and choose the noun-cluster assignment with the lowest error, which we define shortly (Eqn. 6). In practice, this exhaustive search is not a bottleneck as most objects that we are interested in, including those found in the home, contain a manageable number of parts. Note that we also investigated the use of vision-based object recognition to reduce this search space [37], but found the recognition accuracy to be insufficient for such tasks (detectors were prone to false negatives and tend to predict holistic object classes like "bicycle" instead of their parts like "bicycle wheel" and "bicycle frame," which is necessary for our task).

3) Convert Action Verbs to Kinematic Model Type: The final step of our framework converts the parsed action verbs to either "prismatic" or "rotational" kinematic model types.[2] A simple means of performing this conversion, which we treat as an oracle, is to manually create verb dictionaries that span the variety of words that can be used for each of the rotational and prismatic motion types. Figure 3 enumerates the words that define our dictionary. Note that some words are shared by both dictionaries (e.g., "push" can be used to describe both prismatic and rotational motion), in which case the lingual observation would have equal likelihood for different models.

The manual dictionary simply provides an oracle baseline. Our system employs a general, non-manual approach to convert verbs to their corresponding type. Specifically, we embed words in a learned, high-dimensional space and use their relative distances in this space to identify model correspondence. First, we select a small seed dictionary $W = w_1, w_2, ..., w_s$ that includes the s most common words for each model type (we use $s = 3$ in our experiments) from the full manual dictionary in Figure 3. These can be thought of as the seed clusters representing each model type. A seed dictionary is important to construct the model type's centroid vector because there is no single canonical word that can represent the entire meaning of a general kinematic model type such as "prismatic." We use the set $W_{\text{prismatic}} = \{$"shift", "insert", "extract"$\}$ as the seed cluster for the prismatic model and the set $W_{\text{rotational}} = \{$"rotate", "circle", "twist"$\}$ as the seed cluster for the rotational model in our experiments. Next, we convert each word to a d-dimensional word embedding space using word2vec [38], a popular neural language model. We compute the mean over the word vectors in each seed dictionary to arrive at a "centroid" vector $\vec{w}_{\text{prismatic}}$ and $\vec{w}_{\text{rotational}}$ that represents the corresponding kinematic model type. Given a new unseen verb w_{new} from a test sentence, we project it to the same word embedding space (using word2vec) and then compare it with with the centroid vector of each model type according to cosine distance. The model type M_{new} with the smallest distance is set as the model type of this action verb embedding $\vec{w}_{\text{new}}$:

$$M_{\text{new}} = \underset{m \in \{\text{rot,pri}\}}{\arg\min}\ \text{dist}(\vec{w}_m, \vec{w}_{\text{new}}). \tag{3}$$

We require this distance to be lower than that of the other model by a margin (we use 0.1 in our experiments). Otherwise, we treat the word as ambiguous and assign it to both kinematic models.

C. Combining Visual and Lingual Observations

The final step in our framework selects the kinematic graph structure $\hat{\mathcal{M}} = \{\hat{M}_{ij}, \forall (ij) \in E_G\}$ that best explains the visual and lingual observations $D_z = \{D_v, D_l\}$ from the space of all possible kinematic graphs. We do so by maximizing the conditional posterior over the model type associated with each edge in the graph $(ij) \in E_G$:

$$\hat{M}_{ij} = \underset{M_{ij}}{\arg\max}\ p(M_{ij}|D_z) \tag{4a}$$

$$= \underset{M_{ij}}{\arg\max} \int p(M_{ij}, \theta_{ij}|D_z) d\theta_{ij} \tag{4b}$$

Evaluating this likelihood is computationally prohibitive, so we use the Bayesian Information Criterion (BIC) score as an approximation

$$BIC(M_{ij}) = -2 \log p(D_z|M_{ij}, \hat{\theta}_{ij}) + k \log n, \tag{5}$$

[2]We currently assume no action verbs for the "rigid" type and default to the visual observation.

where $\hat{\theta}_{ij}$ is the maximum likelihood parameter estimate (Eqn. 2), k is the number of parameters of the current model and n is the number of visual and lingual observations. We choose the model with the lowest BIC score:

$$\hat{M}_{ij} = \arg\min_{M_{ij}} BIC(M_{ij}) \tag{6}$$

While our previous method [4] only considers visual observations, our new framework performs this optimization over the joint space of visual and lingual observations. Consequently, the BIC score becomes

$$\begin{aligned} BIC(M_{ij}) = -2\Big(&\log p(D_v|M_{ij}, \hat{\theta}_{ij}) \\ &+ \log p(D_l|M_{ij}, \hat{\theta}_{ij}) \Big) + k \log n, \end{aligned} \tag{7}$$

where we have made the assumption that the lingual and visual observations are conditionally independent given the model and parameter estimate. Here, the language model can take one of two forms. The first acts as a hard assignment of a verb to its corresponding model type, whereby we assign a likelihood of one to the model whose centroid vector is closest in the embedding space (Eqn. 3) and zero to the other, subject to a margin. For ambiguous words, i.e., those that can be associated with either model type according to the margin, the probability is equal (0.5) for both of the candidate kinematic models. The second form acts as a soft assignment, setting the model conditional distributions $p(D_l|M_{ij}, \hat{\theta}_{ij})$ according to the cosine similarity between the word in the input associated with the motion $\vec{w}_{\text{verb}} \in D_l$ and the model's centroid vector $\vec{w}_m \in \{\vec{w}_{\text{prismatic}}, \vec{w}_{\text{rotational}}\}$ in the embedding space

$$p(D_l|M_{ij}, \theta_{ij}) = \text{dist}(\vec{w}_m, \vec{w}_{\text{verb}}). \tag{8}$$

We then estimate the overall kinematic structure by solving for the minimum spanning tree of the graph, where we define the cost of each edge as $\text{cost}_{ij} = -\log p(M_{ij}, \theta_{ij}|D_z)$. Such a spanning tree constitutes the kinematic graph that best describes the visual and lingual observations.

IV. Results

We evaluate our framework on 28 RGB-D videos in which a user manipulates a variety of common household and office objects (e.g., a microwave, refrigerator, and drawer). AprilTags [39] were placed on each of the objects parts and used as an observation of ground-truth motion. We mask the AprilTags when running the visual pipeline so as to not affect feature extraction. Of the 28 videos, 13 involve single-part objects and 15 involve multi-part objects. The single-part object videos are used to demonstrate that the addition of lingual observations can only improve the accuracy of the learned kinematic models. The extent of these improvements on single-part objects is limited by the relative ease of inference of single degree-of-freedom motion. In the case of multi-part objects, the larger space of candidate kinematic graphs makes visual-only inference challenging as feature tracking errors may result in erroneous estimates of the graph structure. These experiments are meant to evaluate the extent to which multimodal learning improves model selection.

After watching the videos, we asked a user to provide a single caption for each video. Before doing so, we provided the user with some examples of potential captions that discuss the movement of the individual parts as opposed to single, high-level captions. An example of such a narration is "A man pushes the bicycle frame forward. The front wheel is spinning. The back wheel is rotating." as opposed to the high-level caption "A man pushes a bicycle forward," which would not be sufficient (because our system is unable to associate "bicycle" with only the frame). This is similar to discussions for image and video captioning and question-answering research, where it is well-known that a more detailed, database-like caption is more useful for capturing multiple salient events in the image/video, and for answering questions made of them [34].

A. Evaluation Metrics and Baselines

We estimate the ground-truth kinematic models by performing MAP inference based upon the motion trajectories observed using AprilTags. We denote the resulting kinematic graph as G^*. The kinematic type and parameters for each object part pair are denoted as M^*_{ij} and θ^*_{ij}, respectively. Let $\hat{G}$, $\hat{M}_{ij}$, $\hat{\theta}_{ij}$ be the estimated kinematic graph, kinematic type, and parameters for each object pair from the RGB-D video, respectively.

The first metric that we consider evaluates whether the vision component estimates the correct number of parts. We determine the ground-truth number of parts as the number of AprilTags observed in each video, which we denote as N^*. We indicate the number of parts (motion clusters) identified by the visual pipeline as N_v. We report the average success rate when using only visual observations as $S_v = \frac{1}{K}\sum_{k=1}^{K} \mathbb{1}(N_v^k = N^{k*})$, where K is the number of videos for each object type.

Next, we consider two metrics that assess the ability of each method to estimate a graph with the same kinematic model as the ground truth G^*. The first metric requires that the two graphs have the same structure, i.e., $\hat{M}_{ij} = M^*_{ij}, \forall (ij) \in E_{\hat{G}} = E_{G^*}$. This equivalence requires that vision-only inference yields the correct number of object parts and that the model selection framework selects the correct kinematic edge type for each pair of object parts. We report this "hard" success rate S_h in terms of the fraction of demonstrations for which the model estimate agrees with ground truth. Note that this is bounded from above by fraction for which the vision component estimates the correct number of parts. The second "soft" success rate (denoted by S_s) employs a relaxed requirement whereby we only consider the inter-part relationships identified from vision, i.e., $\hat{M}_{ij} = M^*_{ij}, \forall (ij) \in E_{\hat{G}} \subset E_{G^*}$. In this way, we consider scenarios for which the visual system detects fewer parts than are in the ground-truth model. In our experiments, we found that $\hat{G}$ is a sub-graph of G^*, so we only require that the model type of the edges in this sub-graph agree between both graphs. The metric reports the fraction of total

TABLE I
OVERALL PERFORMANCE OF OUR FRAMEWORK

	Object	N^*	N_v	S_v	Vision-Only S_h	Vision-Only S_s	Our Framework S_h	Our Framework S_s	e_{param}
Single-Part	Door	1	1,1	2/2	2/2	2/2	2/2	2/2	1.86°
	Chair	1	1,1,1	3/3	2/3	2/3	**3/3**	**3/3**	3.34°
	Refrigerator	1	1,1,1,1	4/4	3/4	3/4	**4/4**	**4/4**	5.74°
	Microwave	1	1,1	2/2	2/2	2/2	2/2	2/2	2.02°
Multi-Part	Drawer	2	1,2	1/2	0/2	1/2	**1/2**	**2/2**	0.11°
	Monitor	3	1,1,1,1,3,3	2/6	1/6	4/6	**2/6**	**6/6**	7.27°
	Bicycle	3	1,2,2,2,2,2	0/6	0/6	3/6	0/6	**6/6**	11.33°
	Chair	2	1,2,2	2/3	0/3	1/3	**2/3**	**3/3**	3.05°

demonstrations for which the estimated kinematic graph is a correct sub-graph of the ground-truth kinematic graph.

Once we have the same kinematic models for both $\hat{G}$ and G^*, we can compare the kinematic parameters $\hat{\theta}_{ij}$ to the ground-truth values θ^*_{ij} for each inter-part model $\hat{M}_{ij}$. Note that for the soft metric, we only compare kinematic parameters for edges in the sub-graph, i.e., $\forall(ij) \in E_{\hat{G}} \subset E_{G^*}$. We define the parameter estimation error for a particular part pair as the angle between the two kinematic parameter axes

$$\mathrm{e}_{ij} = \arccos \frac{\hat{\theta}_{ij} \cdot \theta^*_{ij}}{\|\hat{\theta}_{ij}\| \|\theta^*_{ij}\|}, \tag{9}$$

where we use the directional and rotational axes for prismatic and rotational degrees-of-freedom, respectively. We measure the overall parameter estimation error e_{param} for an object as the average parameter estimation error over each edge in the object's kinematic graph. We report this error further averaged over the number of demonstrations.

B. Results and Analysis

Table I summarizes the performance of our multimodal learning method using our embedding-based language model with hard alignment, comparing against the performance of the vision-only baseline [4]. The table indicates the ground-truth number of parts for each object (N^*), a list of the number of parts identified using visual trajectory clustering for each demonstration (N_v), and the fraction of videos for which the correct number of parts was identified (S_v). We then present the hard (S_h) and soft (S_s) model selection rates for our method as well as for the baseline. Our method bests the vision-only baseline in estimating the full kinematic graph for five of the eight objects, matching its performance on the remaining three objects. Specifically, our framework yields accurate estimates of the full kinematic graphs for six more demonstrations than the vision-only baseline, two more for single-part objects and four more for multi-part objects, corresponding to a 21% absolute improvement. Similarly, we are able to estimate a valid sub-graph of the ground-truth kinematic graph for all 28 demonstrations, whereas the vision-only baseline fails to estimate valid sub-graphs for ten of the videos (two for single-part and eight for multi-part objects), corresponding to a 36% absolute improvement.[3] One notable object on which both methods have difficulty is the bicycle for which the trajectory clustering method was unable to identify the presence of the third part (the wheel) due to the sparsity of visual features on the wheels. Consequently, neither method estimated the full kinematic graph for any video, however our framework was able to exploit lingual cues to yield accurate sub-graph estimates for each video. Similarly, clustering failed to identify the three parts that comprise the monitor in all but two videos, for which our method then estimated the correct kinematic graph (and an accurate sub-graph for the remaining four videos).

We then evaluate the accuracy of the parameters estimated by our method by reporting the parameter estimation error for each object, averaged over the set of videos. Note that it is difficult to compare against the error of the vision-only baseline since it does not yield accurate kinematic graphs for several of the videos. When the kinematic graph estimates agree, however, the parameter estimation errors are identical for the two methods, since they both estimate the parameters from the visual data (Eqn. 2).

TABLE II
DETAILED SUCCESS AND FAILURE ANALYSIS

		Manual Dict.		Word Embedding		
		Success	Ambig.	Success	Ambig.	WA
Single-Part	Door	10/10	0/10	10/10	0/10	0/10
	Chair	12/15	3/15	10/15	3/15	2/15
	Fridge	17/20	3/20	17/20	3/20	0/20
	Microwave	10/10	0/10	10/10	0/10	0/10
Multi-Part	Drawer	10/10	0/10	10/10	0/10	2/10
	Monitor	29/30	1/30	29/30	1/30	0/30
	Bicycle	27/30	3/30	19/30	9/30	5/30
	Chair	10/15	5/15	10/15	5/15	0/15

In order to better understand the effects of variability in the linguistic input, we then asked the user to generate

[3]The soft alignment model results in improvements of 14% for the complete graph estimates and 21% for the sub-graphs estimates.

four additional diverse captions for each video. Table II presents the overall performance on the complete set of five captions per video[4] when using our word embedding language model with soft alignment. For comparison, we consider the result with oracle alignment, whereby verbs are matched with their corresponding kinematic type according to the manual dictionary (Fig.3). As can be expected, the model selection accuracy is greater for some captions when using the oracle dictionary. We attribute this difference to two primary factors. First, some captions describe the motion of parts using words that are ambiguous in their meaning. For example, several captions include the term "pull," which may refer to both prismatic or rotational motion according to both the manual dictionary and the word embedding representation, i.e., $\text{dist}(\text{“pull”}, W_{\text{pri}}) \simeq \text{dist}(\text{“pull”}, W_{\text{rot}})$, where $W_{\text{pri}}, W_{\text{rot}}$ are the vectors that represent the two kinematic types. Second, the word embedding-based method may yield inaccurate estimates of word similarity as a result of having been trained on general-domain text. For example, while "slide" is only in the manually defined dictionary for prismatic motion, the word embeddings suggest that it is equidistant from the centroid for each of these types, i.e., $\text{dist}(\text{“slide”}, W_{\text{pri}}) \simeq \text{dist}(\text{“slide”}, W_{\text{rot}})$. Note that ambiguities that result from similarity in the word embedding space distances are different from ambiguities inherent in the verb itself. We attribute the former to the failure of general-domain word embeddings and report this fraction in the last column of Table II, denoted as "WA." The fraction that fail due to the ambiguity inherent in the specific verb itself is denoted as "Ambig." and the fraction that are successful is represented as "Success." The total number of description-video pairs is calculated based up upon five descriptions per video. Note that the multimodal nature of our model allows the visual signal to mitigate ambiguity in the lingual observation. In this way, it is possible to use visual cues to overcome failures of the linguistic models just as we use the lingual signal to mitigate failure of the visual pipeline.

V. Conclusion

We have described a method that uses a joint combination of visual and lingual signals to learn accurate probabilistic models that define the structure and parameters of articulated objects. Our framework treats linguistic descriptions of a demonstrated motion as a complementary observation of the structure of kinematic linkages. We evaluate our framework on a series of RGB-D videos paired with captions of common household and office objects, and demonstrate that the use of lingual cues results in improved model accuracy. Future work includes the incorporation of vision-based object part recognition, using the captions to mitigate noise in the visual recognition. We are also exploring a word embedding representation better suited to this specific domain as means of more efficiently using visual and lingual signals for complex objects. Additionally, we are investigating extending our model to the problem of predicting kinematic models of novel objects, using natural language captions as a means of transferring knowledge from known classes.

[4]We consider a scenario to be successful iff our method identifies the correct kinematic graph for the sub-graph consisting of object parts identified via visual clustering.

References

[1] J. Sturm, C. Stachniss, and W. Burgard, "A probabilistic framework for learning kinematic models of articulated objects," *J. of Artificial Intelligence Research*, vol. 41, no. 2, pp. 477–526, 2011.

[2] X. Huang, I. Walker, and S. Birchfield, "Occlusion-aware reconstruction and manipulation of 3D articulated objects," in *Proc. IEEE Int'l Conf. on Robotics and Automation (ICRA)*, 2012, pp. 1365–1371.

[3] D. Katz, M. Kazemi, J. Andrew Bagnell, and A. Stentz, "Interactive segmentation, tracking, and kinematic modeling of unknown 3D articulated objects," in *Proc. IEEE Int'l Conf. on Robotics and Automation (ICRA)*, 2013, pp. 5003–5010.

[4] S. Pillai, M. R. Walter, and S. Teller, "Learning articulated motions from visual demonstration," in *Proc. Robotics: Science and Systems (RSS)*, Berkeley, CA, July 2014.

[5] V. Ordonez, G. Kulkarni, and T. L. Berg, "Im2Text: Describing images using 1 million captioned photographs," in *Adv. Neural Information Processing Systems (NIPS)*, 2011.

[6] M. Mitchell, X. Han, J. Dodge, A. Mensch, A. Goyal, A. Berg, K. Yamaguchi, T. Berg, K. Stratos, and H. Daumé III, "Midge: Generating image descriptions from computer vision detections," in *Proc. Assoc. for Computational Linguistics (ACL)*, 2012, pp. 747–756.

[7] A. Sharma, A. Kumar, H. Daume III, and D. W. Jacobs, "Generalized multiview analysis: A discriminative latent space," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2160–2167.

[8] C. Kong, D. Lin, M. Bansal, R. Urtasun, and S. Fidler, "What are you talking about? text-to-image coreference," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014.

[9] V. Ramanathan, A. Joulin, P. Liang, and L. Fei-Fei, "Linking people in videos with their names using coreference resolution," in *Proc. European Conf. on Computer Vision (ECCV)*, 2014, pp. 95–110.

[10] S. Guadarrama, E. Rodner, K. Saenko, N. Zhang, R. Farrell, J. Donahue, and T. Darrell, "Open-vocabulary object retrieval," in *Proc. Robotics: Science and Systems (RSS)*, Berkeley, CA, July 2014.

[11] A. Karpathy and L. Fei-Fei, "Deep visual-semantic alignments for generating image descriptions," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[12] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[13] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio, "Show,

attend and tell: Neural image caption generation with visual attention," in *Proc. Int'l Conf. on Machine Learning (ICML)*, 2015.

[14] R. Kiros, R. Salakhutdinov, and R. S. Zemel, "Unifying visual-semantic embeddings with multimodal neural language models," *arXiv:1411.2539*, 2014.

[15] J. Mao, W. Xu, Y. Yang, J. Wang, and A. Yuille, "Deep captioning with multimodal recurrent neural networks (m-RNN)," *arXiv:1412.6632*, 2014.

[16] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, "Long-term recurrent convolutional networks for visual recognition and description," *arXiv:1411.4389*, 2014.

[17] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[18] X. Chen and C. L. Zitnick, "Mind's eye: A recurrent visual representation for image caption generation," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.

[19] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville, "Describing videos by exploiting temporal structure," *arXiv preprint arXiv:1502.08029*, 2015.

[20] N. Srivastava, E. Mansimov, and R. Salakhutdinov, "Unsupervised learning of video representations using LSTMs," *arXiv preprint arXiv:1502.04681*, 2015.

[21] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko, "Sequence to sequence–video to text," *arXiv preprint arXiv:1505.00487*, 2015.

[22] H. Mei, M. Bansal, and M. R. Walter, "Listen, attend, and walk: Neural mapping of navigational instructions to action sequences," in *Proc. Nat'l Conf. on Artificial Intelligence (AAAI)*, Phoenix, AZ, February 2016.

[23] J. Yan and M. Pollefeys, "A general framework for motion segmentation: Independent, articulated, rigid, non-rigid, degenerate and non-degenerate," in *Proc. European Conf. on Computer Vision (ECCV)*, 2006, pp. 94–106.

[24] D. Katz, A. Orthey, and O. Brock, "Interactive perception of articulated objects," in *Proc. Int'l. Symp. on Experimental Robotics (ISER)*, 2010.

[25] K. Hausman, S. Niekum, S. Ostenoski, and G. S. Sukhatme, "Active articulation model estimation through interactive perception," in *Proc. IEEE Int'l Conf. on Robotics and Automation (ICRA)*, Seattle, WA, May 2015.

[26] T. Kollar, S. Tellex, D. Roy, and N. Roy, "Toward understanding natural language directions," in *Proc. ACM/IEEE Int'l. Conf. on Human-Robot Interaction (HRI)*, 2010.

[27] C. Matuszek, D. Fox, and K. Koscher, "Following directions using statistical machine translation," in *Proc. ACM/IEEE Int'l. Conf. on Human-Robot Interaction (HRI)*, 2010.

[28] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. Teller, and N. Roy, "Understanding natural language commands for robotic navigation and mobile manipulation," in *Proc. Nat'l Conf. on Artificial Intelligence (AAAI)*, 2011.

[29] D. L. Chen and R. J. Mooney, "Learning to interpret natural language navigation instructions from observations," in *Proc. Nat'l Conf. on Artificial Intelligence (AAAI)*, 2011.

[30] Y. Artzi and L. Zettlemoyer, "Weakly supervised learning of semantic parsers for mapping instructions to actions," *Trans. Assoc. for Computational Linguistics*, vol. 1, pp. 49–62, 2013.

[31] M. R. Walter, S. Hemachandra, B. Homberg, S. Tellex, and S. Teller, "Learning semantic maps from natural language descriptions," in *Proc. Robotics: Science and Systems (RSS)*, Berlin, Germany, June 2013.

[32] F. Duvallet, M. R. Walter, T. Howard, S. Hemachandra, J. Oh, S. Teller, N. Roy, and A. Stentz, "Inferring maps and behaviors from natural language instructions," in *Proc. Int'l. Symp. on Experimental Robotics (ISER)*, Marrakech/Essaouira, Morocco, June 2014.

[33] S. Hemachandra, F. Duvallet, T. M. Howard, N. Roy, A. Stentz, and M. R. Walter, "Learning models for following natural language directions in unknown environments," in *Proc. IEEE Int'l Conf. on Robotics and Automation (ICRA)*, May 2015.

[34] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "VQA: Visual question answering," *arXiv preprint arXiv:1505.00468*, 2015.

[35] J. Malmaud, J. Huang, V. Rathod, N. Johnston, A. Rabinovich, and K. Murphy, "What's cookin'? Interpreting cooking videos using text, speech and vision," *arXiv preprint arXiv:1503.01558*, 2015.

[36] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky, "The stanford corenlp natural language processing toolkit," in *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2014, pp. 55–60.

[37] R. Girschick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2014.

[38] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Adv. Neural Information Processing Systems (NIPS)*, 2013, pp. 3111–3119.

[39] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *Proc. IEEE Int'l Conf. on Robotics and Automation (ICRA)*, May 2011.